

Patent Application
Docket No. 2003-0018
61922-00006USPT

CUSTOMER NO. 26652

I hereby certify that this correspondence is being deposited with the U.S. Postal Service as Express Mail, Airbill No. EL258053679US in an envelope addressed to: MS Patent Application, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on the date shown below.

Dated: 2/18, 2004 Signature: Carol Marstaller
(Carol Marstaller)

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR U.S. LETTERS PATENT

Title:

**A METHOD FOR DISTRIBUTED DENIAL-OF-SERVICE ATTACK MITIGATION
BY SELECTIVE BLACK-HOLING IN MPLS VPNS**

Inventors:

Thusitha Jayawardena
Luis E. Morales

Steven R. Greenfield
JENKENS & GILCHRIST, A
PROFESSIONAL CORPORATION
1445 Ross Avenue, Suite 3200
Dallas, Texas 75202
(214) 855-4789

**A METHOD FOR DISTRIBUTED DENIAL-OF-SERVICE ATTACK MITIGATION
BY SELECTIVE BLACK-HOLING IN MPLS VPNS**

BACKGROUND OF THE INVENTION

5 Field of the Invention

The present invention relates generally to Virtual Private Networks (VPN) based on Multiprotocol Label Switching (MPLS), and more particularly to redirecting or rerouting VPN traffic in response to an attack caused by an attacker flooding a victim's host system with one or more of several types of attack traffic.

10 Background of the Invention

In a Distributed Denial of Service (DDoS) attack, an attacker takes control of one or more hosts (daemons) and uses the daemons to send an enormous amount of traffic to a, for example, web site so that no other traffic can get through to the website. In essence, the website is clogged or jammed with traffic. The frequency of DDoS attacks in the Internet has
15 grown in the past several years. The flooding of a victim's host system with attack traffic causes legitimate users of the victim's host system to be denied access to applications running on the System-Under-Attack (SUA). The application can be a web server, a file server, a Domain Name System (DNS) server, or other Internet related service or device. The legitimate users cannot access the application due to Central Processing Unit (CPU) and/or
20 bandwidth exhaustion on the system under attack. An attack may have a distributed nature due to the attack traffic being from random, usually spoofed, source IP addresses and originating from many daemon hosts. Also, the attack traffic may enter the victim's Virtual Private Network (VPN) network from various entry points.

10

15

20

3

traffic addressed to the VPN application that is routed through one or more of the edge routers into the ISP network.

In the embodiment of the present invention, the black-hole router is adapted to inject a dummy or bogus IP address into the ISP network. The bogus IP address is the same address
5 as the VPN application's address but has a higher preference value and a community value that selects a number of the edge routers and requires them to address VPN traffic, having the VPN application's IP address, to the bogus address thereby black-holing the traffic. Embodiments of the present invention are well suited for handling DDoS attacks on a website such that some of the traffic is diverted from getting to the website-under-attack and other
10 VPN traffic is allowed to continue to the website-under-attack.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other advantages of the invention will become more apparent upon reading the following Detailed Description of Embodiments of the Invention in conjunction
15 and reference with the drawings.

FIGURE 1 depicts a general ISP network operating in part as a MPLS VPN using a known total black-holing technique;

FIGURE 2 is a graph comparing traffic profiles that may be utilized with an embodiment of the present invention;

20 FIGURE 3 depicts a general ISP network operating in part as a MPLS VPN using a selective black-holing technique in accordance with an embodiment of the present invention; and

FIGURE 4 depicts exemplary black-holing in a MPLS VPN using MP-iBGP route filtering.

While the invention is susceptible to various modifications and alternative forms, (for example, the invention can be easily adapted to a case with multiple black-hole routers.)
5 specific embodiments have been shown by way of example in the drawings and will be described in detail herein. It should be understood, however, that the invention is not intended to be limited to the particular forms disclosed. Rather, the invention is to cover all reasonable modifications, equivalents, and alternatives falling within the spirit and scope of the invention as defined by the appended claims.

10

DETAILED DESCRIPTION OF THE INVENTION

The present invention will now be described more fully hereinafter with reference to the accompanying drawings in which preferred embodiments of the invention are shown. This invention may, however, be embodied in many different forms and should not be
15 constructed as limited to the embodiments set forth herein; rather, these embodiments are provided so that this disclosure will be thorough and complete, and will fully convey the scope of the invention to those skilled in the art.

Referring to FIGURE 1, currently the prevalent black-holing defense against a DDoS attack consists of redirecting all of the traffic destined to the system under attack to another
20 router (or host) called a black-hole router (or host) 12. The redirected attack traffic is then analyzed to determine the origin of the attack or, the redirected traffic is simply dropped into a “black-hole” (hence the term black-holing). In either case, all of the traffic addressed for

the system-under-attack 14 is redirected to the black-hole router 12. The system-under-attack 14 becomes unavailable to all users, legitimate or attacking, for the duration of the black-holing. Total (non-selective) black-holing is therefore not an effective defense against a DDoS attack since the attacker's intention of shutting down a VPN application or creating a denial-of-service to the system-under-attack is not averted.

FIGURE 1 further depicts edge routers R1 through R6. Edge routers are routers on the edge of a VPN or the routers of an ISP that are in direct communication with customers. There is a plurality of other routers within or internal to the ISP network 10 that are not specifically shown called core routers. The routers internal to the ISP network that are part of the VPN are basically transparent routers that are used for backbone routing of traffic throughout the VPN. One of the fundamental differences between a MPLS VPN ISP and a regular ISP network is that each intermediate router of a regular ISP network has to determine the BGP next-hop router for every IP address known through BGP. Conversely, in a MPLS VPN ISP network the internal, intermediate routers are not concerned with BGP next-hop routers, but instead are preconfigured to pass traffic from any edge router R1-R6 to any other edge router R1-R6 through the use of MPLS.

Referring to FIGURE 1, one of the exemplary techniques for mitigating an attack is for an MPLS VPN ISP network in accordance with the present invention to introduce a bogus route for traffic to take to the website. A black-hole router 12 is introduced to the MPLS VPN ISP network 10. The black-hole router 12 contains a bogus address to the website-under-attack 14. The bogus address and the website under attack each have the same address except the bogus address has a higher priority or preference. The higher preference requires

that all traffic going to the website-under-attack's address will be diverted or black-holed to the black-hole router 12. The traffic flow arrows 16 indicate the direction of all traffic from the edge routers R1-R6 addressed to the website-under-attack 14 that is redirected to the black-hole router 12 via the bogus address having a higher priority than the website address.

- 5 The bogus address is the same address as the address for the website-under-attack 14, except the bogus address has a higher priority or preference.

Since all the traffic intended for the website-under-attack is redirected to the black-hole router 12, then even the valid, non-attacking, traffic is black-holed. All traffic intended for the website-under-attack is diverted. The attacker has essentially won the attack and shut
10 down the website-under-attack 14 while the MPLS VPN ISP and their customer(s) figure out what to do.

The MPLS VPN ISP analyzes the redirected traffic being received at the black-hole router 12 in order to trace where the attack traffic is coming from so that a better defense to the attack can be established. Meanwhile, legitimate users of the system are denied access to
15 the website-under-attack 14 because legitimate traffic is also being black-holed.

There are a couple of popular types of attack traffic used by attackers. For example, SYN attack traffic is used as part of a message sent using Transmission Control Protocol (TCP). In TCP, a packet of information that initiates a TCP connection is called a SYN packet. An attacker sends many, many SYN packets to the website-under-attack. In turn, the
20 website-under-attack, responds by attempting to create a TCP connection to all the incoming SYN packets. The proper response to a SYN packet by a website is to reply to each SYN packet with another SYN packet and then wait for another type of packet called an ACK

packet to acknowledge receipt of the website's SYN packet and complete the "TCP Handshake" establishing a TCP connection between the website and the user.

In a SYN attack, the ACK is never sent back to the website-under-attack. Thus, the website-under-attack is bombarded with SYN packets. The website-under-attack replies to each SYN packet with appropriate SYN packets and waits for an ACK for each reply, but
5 never gets any. The website-under-attack receives millions of SYN packets, but never receives an ACK. The website-under-attack is trying to set up all the potential TCP connections, but the TCP handshakes are never completed. The website-under-attack's bandwidth and processor is consumed by servicing the incomplete TCP handshakes and is in
10 effect shut down to valid traffic.

In the other type of ICMP attack called a PING attack, the attacker utilizes the "PING" utility of the ICMP. The PING utility is used to determine whether a specific IP address is accessible. A PING packet is sent to a specified address and a reply is then sent by the specified address. In a PING attack, millions of PINGs are sent to the website-under-
15 attack and the CPU becomes overwhelmed by attempting to respond to all the PINGs. In effect, the bandwidth of both the website's CPU(s) and communication channels is used up so that valid traffic to the website-under-attack is blocked.

When all traffic to a website-under-attack is diverted by an MPLS VPN ISP to a black-hole router 12, it is referred to as non-selective black-holing. In accordance with an
20 embodiment of the invention, a new technique that provides additional options for handling an attack on an application associated with an ISP is called selective-black-holing. As depicted in FIGURE 3, selective-black-holing comprises selectively diverting traffic,

destined to the system-under-attack 14, that enters a MPLS VPN ISP network 10 from a subset of entry nodes (R1, R2, R3) in the MPLS VPN ISP network and allowing traffic entering the MPLS VPN ISP network 10 from other entry points (R4, R5, R6) to reach the system-under-attack 14. This is achieved via dynamic routing protocols in combination with
5 community-based route filtering.

The ISP must either determine or predetermine which edge routers should black-hole potential attack traffic and which edge router can allow the potential attack traffic to proceed to the address of the system-under-attack 14. There are different schemes that can be used to
) choose the entry nodes (routers) or edge routers that implement black-holing of potential
10 attack traffic. One approach is to selectively black-hole traffic consecutively at each entry point (edge router) of the MPLS VPN ISP network. The black-holed traffic can be analyzed to determine the ratio of attack traffic to legitimate traffic at each entry point. Once traffic that is addressed to the system-under-attack from each selected or all entry points is analyzed, black-holing of the traffic can be limited to the entry points that have the highest percent of
15 attack traffic.

Another technique, as shown in FIGURE 2, for using selective-black-holing to analyze potential origins of a DDoS attack is to characterize the intensity of the attack at the major entry points of the DDoS attack traffic by comparing an “average” profile of traffic to a “current” profile of traffic (i.e., a profile from the past five or ten minute interval) at each of
20 the entry points of the MPLS VPN ISP network. The average profile of traffic can be pre-constructed by the service provider for selected, or premium customers of the network. For

example, the program utility called "NETFLOW", Cisco Corporation's traffic statistics collection feature, can be used to construct such profiles.

For example, in FIGURE 2, router one R1 may be an edge router entry point in California. Router two R2 may be an edge router entry point in New York. Router three R3
5 may be an edge router in Boston. Routers R4-R6 may be edge routers in other major cities. By comparing an average traffic profile with the current traffic profile for each router one may surmise that attack traffic mainly originating from edge routers R1, R2 and R3 because the traffic addressed to the system-under-attack 14 increased significantly in the current traffic profile (i.e., during the attack) when compared to the average traffic profile.

10 By using selective-black-holing in accordance with an embodiment of the invention VPN traffic sent to the system under attack via routers one, two and three (R1, R2 and R3) can be, black-holed to the black-hole router 12 shown in FIGURE 3. The traffic received at the black-hole router can be analyzed to determine whether attack traffic is present. And, if so, where the attack traffic is originating.

15 Meanwhile, routers four, five, and six (R4, R5, R6) can selectively be allowed to continue forwarding traffic, addressed to the system-under-attack, to the system-under-attack 14. Thus, the attacker has not succeeded in "shutting down" the system-under-attack 14 because traffic may still be received and transmitted from the system-under-attack 14. The communication bandwidth and/or the microprocessor(s) at the system-under-attack are not
20 attempting to operate beyond their capabilities when majority of the attack traffic is being black-holed. There are various advantages to using selective-black-holing. For, the response time to a DDoS attack can be decreased since various selective-black-holing strategies can be

predetermined and implemented as soon as a DDoS attack is discovered to be underway. By using Border Gateway Protocol (BGP) routing and community-based route filtering in exemplary embodiments of the invention, a predetermined black-holing strategy can be set-up and implemented. BGP is an exterior gateway routing protocol that enables groups of
5 routers to share routing information. BGP is commonly used within and between ISPs. It is understood that an exemplary network or system can utilize other dynamic routing protocols, besides BGP, as long as they allow route filtering. Once an attack has commenced, the selective-black-holing is initiated by communicating the IP address of the system-under-attack with a predefined community value and higher preference from selected point(s) in the
10 ISP network. Depending on the community value, traffic destined to the system-under-attack from any number of entry points (i.e., edge routers) can be rerouted to the black-hole router.

The black-holing strategy can also be adjusted in real-time during an attack. By adjusting the number of entry points (edge routers) that black-hole traffic that is addressed to the system-under-attack, the system-under-attack can be made available to part of the user-
15 community during a DDoS attack.

An embodiment of the present selective-black-holing technique can selectively adjust the number of edge routers (from none to all the routers) that direct a specific type of traffic to a black-hole router in the MPLS VPN ISP network. For route filtering in a selective black-holing MPLS VPN ISP system - iBGP (MP-iBGP) route filtering is utilized using community
20 values. As stated above, selective-black-holing can be implemented using other dynamic routing protocols and filtering schemes besides iBGP.

Still referring to FIGURE 3, black-holing of attack traffic is achieved by injecting, by the ISP (via any router including the black-hole router 12), the IP address of the system-under-attack 14 into a routing protocol, such as a BGP. The injected IP address is the same IP address as the address of the system-under-attack except that the injected IP address has a
5 higher preference than the IP address of the system-under-attack.

In another embodiment of the present invention, route reflectors (RRs) can be implemented. Route reflectors are typically used in large ISPs. Referring to FIGURE 4, the implementation of RRs in an embodiment of the invention comprises using a community attribute in the BGP to enable selective filtering of routing updates in MP-iBGP sessions.
10 Every edge router R1-R6 must communicate with every other edge router in the ISP in order to disseminate information indicating, among other things, the destinations that the edge router can communicate with. In a large network disseminating the information to all the edge routers by each edge router is a daunting, bandwidth consuming task. There can be hundreds to thousands of edge routers in a given ISP network. A technique for simplifying
15 the task of each edge router communicating with all other edge routers is to use route reflectors.

Instead of communicating with every edge router, each edge router informs its immediate route reflector of the possible route(s) it can reach. The route reflector then reflects the route information, via an MP-iBGP mesh 22, to other route reflectors, which in
20 turn communicates to their associated edge routers the routing information.

In FIGURE 4, all entry points (R1-R6) of an ISP Network 40 and the black-hole router 12 are Provider Edge (PE) routers in the VPN context. The PE routers are route

reflector clients of the RRs 16, 18, 20, and 24. Selective route filtering is achieved by filtering based on community values in the MP-iBGP sessions. There is an MP-iBGP full mesh 22 among all RRs 16, 18, 20, 24. In this embodiment, the selectivity is limited to sets of RR clients, edge routers that are served by their own RR. That is, all of the attack and
5 legitimate traffic entering a set of edge routers (e.g., R1 and R2) addressed and destined for the system-under-attack 14 has to be black-holed in the black-hole router 12 or allowed to reach the system-under-attack 14.

Embodiments of the present invention can be adjusted or implemented as a protection for premium customers because average traffic profiles for each edge router can be pre-
10 constructed for the premium customer's IP addresses. When such a customer's application 14 becomes the target of a DDoS attack, the black-hole router 12 can be preconfigured to handle an attack on application 14 by utilizing a set of community values and associated higher preferences for selective black-holing.

Embodiments of the invention provide a real-time adjustable and selectable black-
15 holing strategy for handling a DDoS attack in a manner that does not shut down the system-under-attack. Embodiments of the present invention and obvious variations thereof are contemplated as falling within the spirit and scope of the claimed invention, which is set forth in the following claims: